UNDER REVIEW. NOT FOR DISTRIBUTION

A Bayesian Reliability Analysis of Neutron Induced Errors in High Performance Computing Hardware.[1]

Curtis B. Storlie, Sarah E. Michalak, Heather M. Quinn, Andrew J. DuBois,

Steven A. Wender, David H. DuBois

Los Alamos National Laboratory

Date: January 24, 2012

Abstract

A soft error is an undesired change in an electronic device's state, e.g., a bit flip in computer memory, that does not permanently affect its functionality. In microprocessor systems, neutron-induced soft errors may cause crashes and silent data corruption (SDC). SDC results when a soft error neither gets corrected nor causes a crash, but instead produces a computational result that is incorrect, without the system issuing a warning or error message. Neutron-induced soft errors are of concern for high performance computing platforms that perform scientific computation for two reasons. First, system crashes typically increase the runtimes of scientific calculations, and second SDC may lead to erroneous scientific results. Through accelerated neutron beam testing of hardware in its field configuration, the frequencies of failures (crashes) and of SDCs in hardware from the Roadrunner platform, the first Petaflop supercomputer, are predicted. The impact of key factors on field performance is investigated and estimates of field reliability are provided. Finally, a novel statistical approach for the analysis of interval censored survival data with mixed effects and uncertainty in the interval endpoints, key features of these data, is presented. A detailed description of the algorithm used to fit the proposed model is provided in the online supplementary material.

Keywords: Silent Data Corruption, High Performance Computing, Neutron Beam, Accelerated Testing, Cox Proportional Hazards, Stochastic Search Variable Selection, Gaussian Process, Mixed Effects.

Running title: HPC Reliability Analysis

Corresponding Author: Curtis Storlie, `storlie@lanl.gov`

[1]LA-UR xx-xxxxx

1 Introduction

Cosmic rays are energetic particles that originate in space and our sun and collide with particles as they transit our atmosphere. As these particles interact with the atmosphere they generate a flux of particles that reach the earth's surface. The interaction of these high energy particles with electronics can induce *soft errors* (Ziegler & Lanford 1981, Baumann 2005). A soft error is an undesired change in an electronic device's state, e.g. an unintentional state switch (bit flip) from 0 to 1, or vice versa, that does not permanently impact the device's functionality. At the Earth's surface approximately 95% of the particles capable of causing soft errors are energetic neutrons with the remainder composed of protons and pions (Ziegler 1996).

When they occur in microprocessor systems, neutron-induced soft errors can lead to crashes and silent data corruption (SDC) (Michalak, Harris, Hengartner, Takala & Wender 2005, Constantinescu 2005, Ando, Kan, Tosaka, Takahisa & Hatanaka 2008, Rao, Hong, Sanda, Ackaret, Barrera, Yanez, Mitra, Kellington & McBeth 2008, Sanda, Kellington, Kudva, Kalla, McBeth, Ackaret, Lockwood, Schumann & Jones 2008, Hong, Michalak, Graves, Ackaret & Rao 2009, Michalak, DuBois, Storlie, Rust, DuBois, Modl, Quinn, Manuzzato & Blanchard 2011). SDC is the result of a soft error that neither gets corrected nor causes a crash, but instead produces a computational result that is incorrect, without the system issuing any warnings or error messages. In high performance computing (HPC) platforms used for scientific computation, such errors are of concern since system crashes can increase application runtimes and SDC in scientific applications can result in erroneous scientific conclusions. In this study, accelerated neutron-beam testing is used to estimate the frequencies of SDCs and of failures (crashes) in Los Alamos National Laboratory's (LANL) Roadrunner platform (Koch 2008), the first Petaflop supercomputer (Meuer 2008), as well as the impact of certain factors on these frequencies. There are two main contributions of this work: (i) the results of subjecting

Roadrunner hardware in its field configuration to neutron-beam testing permit estimation of the effect of key factors on the cosmic-ray-neutron-induced error (SDC or failure) rate for the Roadrunner platform and (ii) a novel statistical approach for the analysis of interval censored survival data with mixed effects and uncertainty on the interval endpoints, which is necessary for the analysis of these data, is developed. While some of the results of this study have been presented to the electrical engineering and computer engineering audiences (Michalak, DuBois, Storlie, Rust, DuBois, Modl, Quinn, Manuzzato & Blanchard 2011, Michalak, DuBois, Storlie, Quinn, Rust, DuBois, Modl, Manuzzato & Blanchard 2011), the focus here is on the statistical approach taken to answer the relevant scientific questions.

1.1 Problem Description

This section presents the test methodology and setup, with a more detailed description of the test setup and data collection technology provided in Michalak, DuBois, Storlie, Rust, DuBois, Modl, Quinn, Manuzzato & Blanchard (2011).

Roadrunner includes 17 Connected Units, each with 180 Triblades (Koch 2008) used for computation, so the susceptibility of Triblades to neutron-induced faults is of primary interest. Thus, accelerated testing of Triblades was conducted in October 2009 by running various computational codes on these systems in a neutron *flux* (neutron intensity in neutrons/cm^2/sec) several orders of magnitude higher than that at the earth's surface. The testing was conducted at the Los Alamos Neutron Science Center (LANSCE) Irradiation of Chips and Electronics (ICE) House facility at LANL. Although the neutron flux produced at the ICE house is much more intense, its energy spectrum approximates that at the earth's surface (Takala 2006).

Four Triblades (Koch 2008) were used in this study, but only three of the Triblades are used in this analysis, since one of the Triblades (Triblade 2) was only used during the initial test period used to determine the configuration for the beam testing prior to the

Figure 1: The Test Setup. (a) Triblade 2 housed in the BC-H. (b) A spatial schematic of the test environment with the beam source, the Virtex-IIs, and the BC-H housing the Triblade in the left-most configuration. The wall indicates the point at which the beam enters the ICE House test facility.

(a)

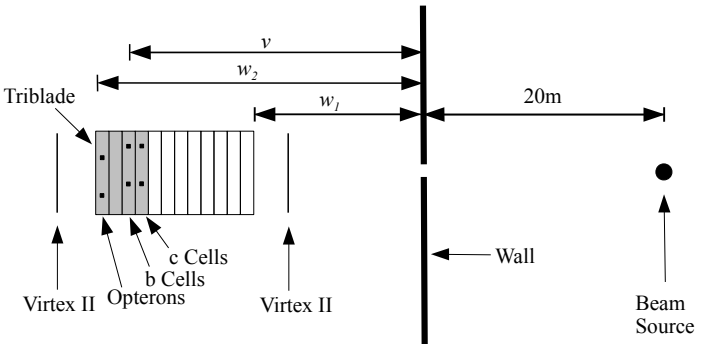

(b)

experimental trials. A Triblade is composed of one IBM LS21 blade with two dual-core AMD Opteron 2210 HE processors, two IBM QS22 blades (QS22a and QS22b) each with two PowerXCell 8i (Cell) processors, and an expansion blade that handles data traffic. For the testing, a Triblade was housed in a BladeCenter-H (BC-H Type 8852) chassis during the experimental trials that involved it (see Figure 1). The BC-H was aligned with the neutron beam passing through the QS22b, the QS22a, the expansion blade, and the LS21, in that order. The test configuration also included two Xilinx Virtex-II FPGAs (Xilinx 2007), one upbeam and the other downbeam of the test system. The Virtex-IIs record a number of bit flips on a memory board, which is commonly treated as a Poisson random variable with a rate proportional to the neutron *fluence* (the accumulated flux over time in neutrons/cm^2) at the Virtex-II (see Section 2.2). These measurements are used for calculating corrected neutron fluence exposures for the hardware under test as explained in Section 2.

Because one study goal was to estimate field susceptibility to neutron-induced faults, the hardware was tested in the field configuration described above and while running different applications, including some used for scientific computation. The test applica-

tions used for the Cell were five computational test codes (cg, corr, hpl, int_add, vpic) detailed in Michalak, DuBois, Storlie, Rust, DuBois, Modl, Quinn, Manuzzato & Blanchard (2011), along with an idle test code, and a varied application condition, which cycled through various applications. The Opteron test applications were an Opteron-only version of the correlator test code (corr) and an idle condition. The test codes were designed to run a job that lasted roughly one minute, produce an output line (start and stop times, the relevant application, the hardware under test, and related output including that needed to determine if an SDC had been observed), and then repeat the cycle until a crash, hang, or an operator decision to terminate the trial. Virtex-II measurements related to neutron fluence exposure were then collected, and the system was rebooted before beginning the next test. The trials were conducted with two different beam widths: the first 52 trials used a 2-inch beam width, while the last 61 used a 1-inch beam width, for a total of $n = 113$ experimental trials.

For a particular trial, a single processor (Cell or Opteron) ran the desired application while all of the hardware in that processor's beampath was irradiated by the beam. That is, a cylindrical volume throughout the Triblade (right to left in Figure 1) under test was exposed to the beam. Therefore, it is not possible to attribute the cause of the error to any one component for certain. Most notably, since two QS22s (a and b) are in a Triblade, when a Cell in one QS22 is running an application, the corresponding Cell in the other QS22 (running the idle test code) was also in the beampath. It is possible (although not that likely) that observed errors result from hardware other than the processors in the beampath, hardware that was affected by scatter, or causes unrelated to the beam.

A key goal of this work is to assess the impact on neutron fluence until error (SDC or failure) of several of the factors mentioned above, namely (i) Application (cg, corr, hpl, int_add, vpic, varied, idle), (ii) Triblade (1, 3, 4), (iii) Beam Aim (Opteron, Cell), and (iv) Beam Width (1-inch, 2-inch). A related goal is prediction of the distribution of failures and SDCs due to cosmic ray induced neutrons, in actual clock time (for the

neutron flux at Los Alamos) on a per beam aim and per Triblade basis, as well as for a Roadrunner Connected Unit (CU) that includes 180 Triblades.

The Cox Proportional Hazards Model (CPHM) framework (Cox 1972) is adopted for this analysis. However, there are several challenges that make the application of the CPHM less than straightforward in our case. The most difficult challenge is that the data are interval censored, with uncertain endpoints. This is because the neutron fluence to which a device was exposed before an error occurs is unknown. Instead, all that is known are the time of the last output line (a_i) and the time at which the operator noticed the error (b_i) (or in some cases, the time at which the operator decided to end the trial before observing an error). To make matters more complicated, the exact neutron fluence to which a particular component in the BC-H was exposed until times a_i or b_i is not known without error. This is because of attenuation (loss in beam intensity when traveling through a medium) of the neutron beam once it enters the Triblade. Since beam attenuation through a Triblade is unknown, the attenuation effect must be treated as a parameter that must be estimated, along with the other parameters in the model, thus leading to the uncertain endpoints a_i and b_i.

In addition to the interval censoring issues, the model must include mixed effects, since the effect of the Triblade under test is most appropriately treated as a random effect. Further, making predictions about the reliability of components in a particular beam aim, Triblades, and Roadrunner CUs requires that the hazard function be explicitly modeled. Finally, it is desirable to perform model selection to aid in prediction and assessing variable importance. To our knowledge, there is currently no methodology available that addresses all of these challenges. The next subsection reviews some of the relevant literature which was built upon to create a data analysis methodology appropriate for this problem.

1.2 Review of Existing Methodology

Finkelstein (1986) was the first to directly address the issue of interval censored data in the CPHM by jointly estimating the cumulative hazard and the regression coefficients. Estimation is based on a full likelihood under the proportional hazards model where the cumulative hazard is treated as a step-function at the event times. The step-function assumption is not well suited for the prediction of future events. Further, the number of parameters increases with the number of observations, which causes numerical instability unless some form of penalization is introduced (So, Johnston & Kim 2010).

Goggins, Finkelstein, Schoenfeld & Zaslavsky (1998) avoid the estimation of the hazard function by treating the interval censoring as an incomplete data problem and use an EM algorithm to obtain estimates using just the partial likelihood. Goetghebeur & Ryan (2000) also use EM while nonparametrically modeling the hazard, but assuming a step-function for the hazard. Kooperberg & Clarkson (1996) adopt a full likelihood approach as well, which allows for interval censoring, with linear or natural cubic splines for the log-hazard. Abrahamowicz, Ciampi & Ramsay (1992) also use cubic splines to estimate the log-hazard. These approaches are closer to what is needed here, but do not easily allow for the mixed effects and uncertainty of the censored interval endpoints needed for this analysis.

Tibshirani (1997) performs variable selection in the CPHM via L1-penalization (i.e., LASSO) of the partial log-likelihood. Fan & Li (2002) and Zhang & Lu (2007) achieve penalization of the partial log-likelihood using the smoothly clipped absolute deviation (SCAD) (Fan & Li 2001) and adaptive LASSO penalties, respectively, both of which possess the attractive oracle property. However, none of these methods are directly applicable to interval censored data.

Raftery, Madigan & Volinsky (1995) and Volinsky, Madigan, Raftery & Kronmal (1997) use Bayesian model averaging (BMA) on a reduced subset of *likely* models and

demonstrate a substantial *predictive* improvement over the classical approach of selecting one model. Faraggi & Simon (1998) extends the Bayesian variable selection concept to interval censored data. However, these approaches use the partial likelihood and thus are not Bayesian approaches in the proper sense, nor do they provide an estimate of the hazard function which is needed here in order to make predictions.

Ibrahim, Chen & MacEachern (1999) extends the previous work, using the full likelihood by assuming a discrete Gamma process on the hazard function, thus allowing for interval censoring. Unfortunately, this results in an increasing hazard almost surely, which may not be a reasonable assumption when considering many devices can exhibit a bathtub curve hazard (i.e., a period of a higher failure rate for new devices that decreases as the device ages and possibly increases later).

In the Bayesian paradigm, recent variable selection for models has focused on the Stochastic Search Variable Selection (SSVS) (George & McCulloch 1993, George & McCulloch 1997). In this procedure, the model (i.e., which covariates have a non-zero coefficient) is treated as a random variable and is sampled along with the other unknown parameters in the MCMC routine. Kinney & Dunson (2007) extended SSVS for mixed models in linear and logistic regression. For a good review of Bayesian variable selection methods, see O'Hara & Sillanpää (2009).

Most recently, Lee, Chakraborty & Sun (2011) has applied SSVS to the Cox proportional hazards model using a fully Bayesian approach. This approach is the most related to the approach taken here for the beam testing data. However, for simplicity this work assumes a discrete response model, with a cumulative hazard that is assumed to be a discrete Gamma process. This model for the cumulative hazard function implies that events occur at discrete times, which is not sufficient for predicting very small error times (e.g., a minimum of many error times such as when modeling the time until error for a system composed of many individual components as in Section 4.5). In addition, their approach does not directly address mixed effects and interval censoring, particularly

when the interval endpoints are uncertain, as is needed in this analysis.

In this paper, concepts from many of the works above are used to create an appropriate model for the beam testing data. Namely, SSVS is applied to the CPHM with mixed effects and interval censoring with uncertain endpoints. A fully Bayesian approach is taken by considering the full likelihood and modeling the log-hazard function as a Gaussian Process (GP) (Stein 1999). The particular GP used in this work is a Bayesian parallel to the smoothing spline (Wahba 1990), so this choice is partially motivated by Abrahamowicz et al. (1992) and Kooperberg & Clarkson (1996) who used cubic splines to model the log-hazard. The model for the hazard function used here has the attractive feature of being a function with a specified level of continuity. It is also discussed in Section 4.4 that the Cox regression coefficients are not sensitive to the choice of priors on the parameters for the GP that governs the hazard.

Section 2 presents the Bayesian model used in this analysis along with prior specifications and their rational. Section 3 describes the MCMC sampling routine used to fit this model. Section 4 discusses the results of the fitted model along with verification of assumptions and assessment of prior sensitivity, and Section 5 concludes the paper.

2 Modeling the Experimental Data

Define T_i, as the time until an error (SDC or failure) for each experimental trial $i = 1, \ldots, n$. Most of the T_i are interval censored. That is, an interval (a_i, b_i), such that $T_i \in (a_i, b_i)$, with $0 \leq a_i \leq b_i \leq \infty$ is observed (except in three trials where the operator noticed the failure during the same minute that the last output line was recorded, so that a_i and b_i are the *same* within measurement precision). However, the desire is to model Y_i, $i = 1, \ldots, n$, which is the amount of exposure (or fluence in neutrons/cm^2) to the component under test during the i-th trial until an error occurs. The relationship between Y_i and T_i is non-trivial, but can be approximated based on some physical principles, with Section 2.2 providing details. First the model for a known Y_i (ignoring for now the issues

involved with unknown interval censored endpoints for Y_i) is presented in Section 2.1. While in Section 2 fluence until failure and SDC are modeled collectively as the fluence accumulated until any error, Section 4.5 describes an extension that separates out the events of failure and SDC.

2.1 Model for the Exposure until Error

The Y_i are modeled as independent observations from a CPHM,

$$h_i(y \mid \boldsymbol{x}_i, \boldsymbol{z}_i) = h_0(y) \exp\{\boldsymbol{\beta}'\boldsymbol{x}_i + \boldsymbol{\gamma}'\boldsymbol{z}_i\}, \tag{1}$$

where (i) h_0 is the baseline hazard function, (ii) $\boldsymbol{x}_i = (x_{i,1}, \ldots, x_{i,p})'$ is a vector of covariates corresponding to the fixed effects for the the i-th observation (the test application, beam aim (Cell or Opteron), and beam width (1-inch or 2-inch), which are explained further in the beginning of Section 4), (iii) $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_p)'$ is a vector of fixed effect parameters, (iv) $\boldsymbol{z}_i = (z_{i,1}, \ldots, z_{i,p})'$ is a vector of covariates corresponding to the random effect due to Triblade for the i-th trial, and (v) $\boldsymbol{\gamma} = (\gamma_1, \ldots, \gamma_q)' \sim N(\mathbf{0}, \boldsymbol{\Lambda})$ is a vector of random effect parameters.

The probabilistic model for Y_i is then

$$S_i(y) = P(Y_i > y) = \exp\left\{-\int_0^y h_i(u \mid \boldsymbol{x}_i, \boldsymbol{z}_i)du\right\} \tag{2}$$

where (1) provides the functional form of h_i.

In the general framework presented here, the elements of $\boldsymbol{\beta}$ may be grouped in order to facilitate variable selection. For example, suppose the first m elements of $\boldsymbol{\beta}$ correspond to indicator variables describing the levels of a single factor (e.g., which application is running). In this case, it may be desirable to have a single indicator variable to permit variable selection on the entire factor or group of βs. Therefore, the modeling includes $g_f \leq p$ groups denoted by $J_1, \ldots, J_{g_f}$, where J_l is a set of indices such that $\{\beta_j : j \in J_l\}$ is the set of βs corresponding to the l-th fixed effect group (i.e., factor). A similar construction is used for the random effects $\boldsymbol{\gamma}$. Let the $g_r \leq q$ groups for random effects

be denoted $K_1, \ldots, K_{g_r}$, where K_l is a set of indices such that $\{\gamma_k : k \in K_l\}$ is the set of γs corresponding to the l-th random effect group (i.e., factor).

The following prior distribution is assumed for $\boldsymbol{\beta}$,

$$\begin{aligned} \beta_j &= \delta_l \xi_j, \text{ for } j \in J_l,\, l = 1, \ldots, g_f, \\ \delta_l &\overset{ind}{\sim} \text{Bernoulli}(\dot{a}_l), \\ \xi_j &\overset{ind}{\sim} N(0, \dot{b}_j^2). \end{aligned} \tag{3}$$

The δ_l are indicator variables which permit variable selection for the l-th fixed effect factor (i.e., if $\delta_l = 0$ then all $\beta_j = 0$ for the l-th factor). If $\delta_l = 1$, then the β_j for the l-th factor have independent normal distributions.

Recall that for the random effects, it is assumed that $\boldsymbol{\gamma} \sim N(\mathbf{0}, \boldsymbol{\Lambda})$. Kinney & Dunson (2007) describe a general hierarchical modeling framework for $\boldsymbol{\Lambda}$. However, in many cases, such as that here (i.e., with Triblade as a random effect) it is appropriate to simplify this structure and assume (conditionally) independent random effects, i.e., let $\boldsymbol{\Lambda}$ be a positive diagonal matrix with diagonal elements $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_q)'$, and so that $\gamma_k \overset{ind}{\sim} N(0, \lambda_k)$.

The following prior distribution for $\boldsymbol{\lambda}$ is assumed

$$\begin{aligned} \sqrt{\lambda_k} &= \zeta_l \psi_l, \text{ for } k \in K_l,\, l = 1, \ldots, g_r \\ \zeta_l &\overset{ind}{\sim} \text{Bernoulli}(\dot{c}_l), \\ \psi_l &\overset{ind}{\sim} \text{HC}(\dot{d}_l), \end{aligned} \tag{4}$$

where HC(d) is the half-Cauchy distribution with median d. If $\zeta_l = 0$, then all of the $\lambda_k = 0$ and all of the $\gamma_k = 0$ for the l-th random effect factor, while if $\zeta_l = 1$ then γ_k is normal with mean 0 and standard deviation ψ_l which is drawn from a HC distribution. The HC was recommended by Gelman (2006) as a prior for standard deviation parameters, and provided good performance when used in Reich, Storlie & Bondell (2009) as the prior for a random effects standard deviation (with variable selection similar to that here) in the context of nonparametric regression.

For the baseline hazard function h_0 a GP is assumed with covariance

$$K(s,t) = \sum_{m=0}^{2} \dot{f}_m^2 B_m(s) B_m(t) + \tau^2 B_4(|s-t|), \tag{5}$$

where B_m is the m-th Bernoulli polynomial and $\dot{f}_0^2$, $\dot{f}_1^2$, $\dot{f}_2^2$, and τ^2 are variance parameters. This choice of covariance function is also described in Reich et al. (2009) and Wahba (1990). The covariance K is only defined on $[0,1]^2$, so the domain is transformed to $\tilde{t} = t/t^*$, for some large value t^* (t and t^* are measured in neutrons/cm^2, not time). It is further assumed that the baseline hazard $h_0(t) = h_0(t^*)$, a constant, for any $t > t^*$. This is just a formality for predictive purposes, as t^* is chosen so that $t < t^*$ always holds in the likelihood evaluations for estimation purposes. The *constant hazard eventually* assumption also ensures a regularity of the model, in the sense that $\int_0^\infty h_0(t)dt = \infty$ almost surely, so that $S_i(\infty) = 0$ in (2), and Y_i is a proper random variable.

Reich et al. (2009) describe how the mean zero GP with covariance K can be thought of as a quadratic function (with variances on the quadratic coefficients given by $\dot{f}_0^2$, $\dot{f}_1^2$, and $\dot{f}_2^2$, respectively) plus a remaining stationary function (with variance τ^2) that allows for deviation from a quadratic curve. Thus, the formal model for h_0 is

$$\log h_0(t) = \begin{cases} \sum_{m=0}^{2} \phi_m B_m(\tilde{t}) + g_0(\tilde{t}) & \text{for } t \le t^* \\ \log h_0(t^*) & \text{for } t > t^* \end{cases} \tag{6}$$

with $\tilde{t} = t/t^*$, $\phi_m \sim N(\dot{e}_m, \dot{f}_m^2)$, $g_0(t) \sim \text{GP}(0, \tau^2 K^*)$, $K^*(s,t) = B_4(|s-t|)$, and $\text{GP}(0, K)$ is a GP with mean 0 and covariance function K. As in Reich et al. (2009) we assume $\tau \sim \text{HC}(\dot{g})$.

The covariance K was chosen for this application over the more traditional powered exponential covariance because K is more stable when many points are close together, and it leads to a large computational savings as described in Section 3. The covariance function K also has a close ties to the cubic smoothing spline estimator (Wahba 1990).

Values for $\{\dot{a}_l\}_{l=1}^{g_f}$, $\{\dot{b}_j\}_{j=1}^{p}$, $\{\dot{c}_l\}_{l=1}^{g_r}$, $\{\dot{d}_l\}_{l=1}^{g_r}$, $\{\dot{e}_m\}_{m=0}^{2}$, $\{\dot{f}_m\}_{m=0}^{2}$, $\dot{g}$, and t^* need to be chosen to complete the model specification. Section 4.1 describes choices of these values based on previous information and expert solicitation, and Section 4.4 presents a

summary of the effect of these choices on the analysis.

2.2 Exposure to the Component Under Test

In Section 2.1, a model was described for Y_i, $i = 1, \ldots, n$, which is the amount of *exposure* to the component under test (in neutrons/cm^2) until an error (SDC or failure). As mentioned, the Y_i are not directly available in the experimental data. The following describes the relationship between the directly observable quantities a_i and b_i, the interval censored values for the *time* until the i-th error T_i, and Y_i.

Let $\mathcal{E}_i(0)$ denote the total exposure at the wall in neutrons/cm^2 where the beam enters the ICE House facility (see Figure 1) from the start of the i-th trial until time of error T_i. For a known T_i, with negligible error the corresponding $\mathcal{E}_i(0)$ can be calculated based on a known proportionality relation (i.e., the beam flux in neutrons/cm^2/sec at the wall is known). However, the experimental data do not contain the T_i, but rather the interval censored time values (a_i, b_i) for each trial, $i = 1, \ldots, n$. Therefore, $\mathcal{E}_i(0)$ itself is unknown, but the interval censored exposure values *at the wall* for the i-th trial $(\mathcal{A}_i(0), \mathcal{B}_i(0))$, corresponding to (a_i, b_i) can be calculated with negligible error.

Denote the distance from the wall to the component being tested in the i-th trial as v_i. The data needed to fit the model described in Section 2.1 are corresponding intervals for the exposure at the location of the component under test in the Triblade $(\mathcal{A}_i(v_i), \mathcal{B}_i(v_i))$. The precise values of $(\mathcal{A}_i(v_i), \mathcal{B}_i(v_i))$ are *not* known, but will be less than $(\mathcal{A}_i(0), \mathcal{B}_i(0))$ due to *divergence* (the spreading out of the neutron stream so there are fewer accumulated neutrons/cm^2 within the stream further from the beam source). *Attenuation* (i.e., the gradual loss in intensity through a medium like a Triblade) of the beam further affects the exposure to the components in different locations in the interior of the Triblade. The reduction in fluence due to divergence can be calculated under the commonly used assumption that the neutron beam is a point source. With the wall 20 m from the beam source, assume a hypothetical object that is v m from the wall (but in

front of the Triblade so that attenuation effects can be ignored). If $\mathcal{E}_i(0)$ is the exposure at the wall for the i^{th} trial, then the exposure an object v m away from the wall would be

$$\mathcal{E}_i(v) = \mathcal{E}_i(0)\frac{20^2}{(20+v)^2}. \tag{7}$$

The decrease in radiation due to attenuation through variable matter (i.e., the Triblade) is difficult, if not impossible to account for precisely. However, the total proportion reduction through the entire Triblade for each of the four beam aims, $r = 1, 2, 3, 4$ (lower Cell, upper Cell, lower Opteron, upper Opteron, respectively, depcited by the dots in Figure 2(b) (b)), can be estimated with the Virtex-II readings. We allow the reduction due to attenuation (from the front of the front Virtex II to the back Virtex II), α_r, $r = 1, 2, 3, 4$ to be different for each of the four beam aims, because the Triblade has different material in the path of the beam in each blade for each of these beam aims.

The data provide information on the total attenuation from the front Virtex-II to the back Virtex-II, but not on the attenuation at the point of a particular component under test or the locations where the attenuation occurs. However, this information can be used to bound the attenuation at a particular location in the Triblade. The actual exposure to a component during a given experiment is already treated as a censored interval (because the exact clock time of the failure is unknown), so these bounds on the attenuation simply add to the uncertainty in the censored interval.

The model for attenuation widens the censored interval for exposure using the following logic. The lowest the exposure to a particular component in the beam could be is that resulting when all of the reduction due to attenuation happened upbeam of the component. On the other hand, the highest this exposure could be is that resulting when none of the reduction due to attenuation happened upbeam of the component (i.e., all attenuation happens downbeam of the component). Assuming ignorance about anything in between these two extremes results in a second layer of interval censoring for the exposure to the component under test.

Specifically, the expression for exposure to a hypothetical object in the path of the beam for beam aim $r \in \{1, 2, 3, 4\}$ a distance v from the wall is,

$$\mathcal{E}_i(v) = \begin{cases} \mathcal{E}_i(0)\frac{20^2}{(20+v)^2} & \text{for } v \leq w_1 \\ \mathcal{E}_i(0)\frac{20^2}{(20+v)^2}\alpha_r & \text{for } v > w_2 \\ E^* \text{ for some } E^* \in \left[\mathcal{E}_i(0)\frac{20^2}{(20+v)^2}\alpha_r, \mathcal{E}_i(0)\frac{20^2}{(20+v)^2}\right] & \text{for } w_1 < v \leq w_2 \end{cases}, \quad (8)$$

where w_1 and w_2 are the distances from the wall to the front and back of the Triblade BC-H, respectively, as seen in Figure 1, and α_r is the total proportion reduction due to attenuation for the r-th beam aim. If $w_1 < v \leq w_2$ we simply assume ignorance about the value of $\mathcal{E}_i(v)$ other than the fact that it is in an interval.

Finally, the accumulated fluence to the component under test during the i-th trial, $\mathcal{E}_i(v_i)$, is assumed to lie in the interval (a_i^*, b_i^*), where

$$a_i^* = \mathcal{A}_i(0)\frac{20^2}{(20+v_i)^2}\alpha_{r_i} \quad (9)$$

$$b_i^* = \mathcal{B}_i(0)\frac{20^2}{(20+v_i)^2}, \quad (10)$$

where r_i is the beam aim (1,2,3,4) for the i-th trial. The censored interval in (9) and (10) thus accounts for both the incomplete knowledge of the precise *clock time* at which the failure occurred (through $\mathcal{A}_i(0)$ and $\mathcal{B}_i(0)$), and incomplete knowledge of the *attenuation* effect inside the Triblade BC-H (through the inclusion of α_{r_i} or not).

If the proportion reduction due to attenuation for each beam aim α_r, $r = 1, 2, 3, 4$ were known, then the censored interval for the total amount of exposure to the component under test (a_i^*, b_i^*) could be calculated using (9) and (10). However, the rates α_r are unknown parameters that must be estimated, which is what makes the values of these interval censored endpoints (for the total amount of exposure during the i-th trial to the component being tested) unknown. The estimation of the α_r is feasible though, since there is information about the α_r contained in the readings from the front and back Virtex-IIs.

Let $v_{f,i}$ and $N_{f,i}$ be the distance from the wall and the number of bit flips, respectively, for the front Virtex-II, and similarly define $v_{b,i}$ and $N_{b,i}$ for the back Virtex-II for trial i. The number of bit flips on a Virtex-II is known to be approximately Poisson (due to the Poisson approximation to the Binomial for large n, small p) with mean proportional to the neutron fluence times the sensitive area (the cumulative area of the bits on the Virtex II). Therefore, we have

$$\begin{aligned} N_{f,i} &\sim \text{Poisson}(\mu_{f,i}) \\ N_{b,i} &\sim \text{Poisson}(\mu_{b,i}), \end{aligned}$$

where

$$\begin{aligned} \mu_{f,i} &= \kappa \mathcal{C}_i(v_{f,i}) = \kappa \mathcal{C}_i(0) \frac{20^2}{(20+v_{f,i})^2} \\ \mu_{b,i} &= \kappa \mathcal{C}_i(v_{b,i}) = \kappa \mathcal{C}_i(0) \frac{20^2}{(20+v_{b,i})^2} \alpha_{r_i}, \end{aligned} \tag{11}$$

where $\mathcal{C}_i(0)$ is the total fluence at the wall for the entire duration of the ith trial (i.e., usually $\mathcal{C}_i(0) = \mathcal{B}_i(0)$, unless the trial ended with no error occurring, in which case $\mathcal{B}_i(0) = \infty$ and $\mathcal{C}_i(0) = \mathcal{A}_i(0)$), and κ is the proportionality constant (the sensitive area) which is the same for both $\mu_{f,i}$ and $\mu_{b,i}$.

For the reasons discussed previously, the values of $\mathcal{C}_i(0)$ are known (with negligible error), based on the known beam flux at the wall and the time spent in the beam. However, α and κ are (nuisance) parameters that must be estimated. The sensitive area is known to be 3.586×10^{-8} cm^2, but κ is allowed to vary about this known area in the prior specification in order to allow for error in the approximation given in (11); see Section 4.1.

We assume

$$\kappa \sim \text{Gamma}(\dot{h}_1, \dot{h}_2),$$

which is fairly insensitive to $\dot{h}_1, \dot{h}_2$ since there is information about κ in every experimental trial. However, the problem with assuming independent priors for the α_r is that there is

only information about α_r for trials where $r_i = r$. Some beam aims (specifically those for the Opterons) have fewer trials, which would lead to a large uncertainty about these α_r. However, LANL experts suggested that the proportion reduction due to attenuation for the four beam aims should be similar; see Section 4.1. Therefore, we assume that

$$\alpha_r \sim \text{Gamma}(\mu_\alpha, \sigma_\alpha^2),$$

where μ_α and σ_α^2 are the mean and variance, respectively, of the Gamma distribution, and

$$\begin{aligned} \mu_\alpha &\sim \text{Gamma}(\dot{i}_1, \dot{i}_2), \\ \sigma_\alpha^2 &\sim \text{Gamma}(\dot{j}_1, \dot{j}_2). \end{aligned} \tag{12}$$

3 MCMC Routine

The observed data include $\boldsymbol{a} = (a_1, \ldots, a_n)'$, $\boldsymbol{b} = (b_1, \ldots, b_n)'$, $\boldsymbol{N}_f = (N_{f,1}, \ldots, N_{f,n})'$, $\boldsymbol{N}_b = (N_{b,1}, \ldots, N_{b,n})'$, where the a_i and b_i are the interval censored clock times as defined in Section 1.1. In addition, there are the fixed, known covariates, $\boldsymbol{x}_i$, $\boldsymbol{z}_i$, v_i, $v_{f,i}$, and $v_{b,i}$, $i = 1, \ldots, n$, which are suppressed in the likelihood notation below for convenience. As mentioned, most of the observations are interval censored, but some observations are not technically interval censored (i.e., $a_i = b_i$ for three trials in which the operator noticed a failure immediately after it happened), and these observations must be handled differently in the likelihood.

We adopt the *Gelfand* style for density notation (Gelfand 1990), and let $[X]$ denote the probability density function of the random variable X, and $[X \mid Y]$ denote the conditional density of X given Y. The likelihood for the data is then

$$\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \boldsymbol{\phi}, g_0, \tau, \alpha, \kappa)$$

$$
\begin{aligned}
&= [\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{N}_f, \boldsymbol{N}_b \mid \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \boldsymbol{\phi}, g_0, \tau, \alpha, \kappa] \\
&= [\boldsymbol{a}, \boldsymbol{b} \mid \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \boldsymbol{\phi}, g_0, \tau, \alpha] \cdot [\boldsymbol{N}_f \mid \alpha, \kappa, \boldsymbol{a}, \boldsymbol{b}] \cdot [\boldsymbol{N}_b \mid \alpha, \kappa, \boldsymbol{a}, \boldsymbol{b}] \\
&= \prod_{i=1}^{n} \left\{ \left[(S_i(a_i) - S_i(b_i))\, I_{\{a_i < b_i\}} + d_i(a_i) I_{\{a_i = b_i\}} \right] P_{\mu_{f,i}}(N_{f,i}) P_{\mu_{b,i}}(N_{b,i}) \right\}, \quad (13)
\end{aligned}
$$

where S_i is given in (2), $d_i(y) = -\partial S_i(y)/\partial y$, and $P_\mu(\cdot)$ is the density of a Poisson random variable with mean μ.

Inference can then be carried out with standard MCMC techniques; more details are given in the online Supplementary Material. A slight complication is that the calculation of (13) requires evaluation of $\int_0^y g_0(t)dt$ for $y \in \{a_1, b_1, \ldots, a_n, b_n\}$ (recall g_0 is the GP representing deviation from quadratic in $\log h_0$, see (6)). To make this integral evaluation more tractable, a dense grid of points $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_N)'$, equally spaced by Δ_ω, is used to evaluate the integral with a quadrature approximation, where $N = 500$ in this application. This being the case, g_0 need only be known at the grid points $\boldsymbol{\omega}$, so only the values of a realization of g_0 in the MCMC algorithm at $g_0(\boldsymbol{\omega}) = (g_0(\omega_1), \ldots, g_0(\omega_N))'$ are required. For practical purposes then, samples are obtained from the approximate posterior distribution

$$
\begin{aligned}
&[\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \boldsymbol{\phi}, g_0(\boldsymbol{\omega}), \tau, \alpha, \kappa \mid \boldsymbol{a}, \boldsymbol{b}, \boldsymbol{N}_f, \boldsymbol{N}_b] \\
&\qquad \propto \mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \boldsymbol{\phi}, g_0(\boldsymbol{\omega}), \tau, \alpha, \kappa) \cdot [\boldsymbol{\beta}] \cdot [\boldsymbol{\gamma} \mid \boldsymbol{\lambda}] \cdot [\boldsymbol{\lambda}] \cdot [\boldsymbol{\phi}] \cdot [g_0(\boldsymbol{\omega}) \mid \tau] \cdot [\tau] \cdot [\alpha] \cdot [\kappa] \quad (14)
\end{aligned}
$$

where $\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\lambda}, \boldsymbol{\phi}, g_0, \tau, \alpha, \kappa)$ is the same as that in (13) only it is evaluated using the approximation $\int_x^y g_0(t)dt \approx \sum_{\omega_k : x \le \omega_k < y} g_0(\omega_k)\Delta_\omega$, and $[g_0(\boldsymbol{\omega}) \mid \tau]$ is a multivariate normal density with mean $\mathbf{0}$ and covariance $\tau^2\boldsymbol{\Sigma}$ with (k, l)-th element $\tau^2 K^*(\omega_k, \omega_l)$.

The sampling is carried out with a typical hybrid Gibbs updating scheme with several Metropolis Hastings (MH) steps. The details are provided in the online Supplementary Material. In this application, the MCMC algorithm is fairly robust to the precise values of the control parameters.

The algorithm requires the evaluation of the data likelihood and the random generation of $g_0^*(\boldsymbol{\omega})$, which requires the $O(N^3)$ Cholesky decomposition of $\tau^2\boldsymbol{\Sigma}$. However,

the structure of this decomposition is affected only by a constant as τ changes. Since $\boldsymbol{\Sigma}$ remains constant throughout the sampling scheme, the Cholesky decomposition is only needed once outside of the MCMC iterations. This speeds up computation considerably over other popular GP covariances such as the powered exponential, for example, while still maintaining a large amount of flexibility for modeling the baseline hazard. Further, this covariance (being quadratic as opposed to exponential squared) is much more stable computationally (i.e., remains computationally positive definite) when there are many points close together as is the case here.

The estimation procedure was tested on several known truth test cases with interval censored data, but without the complication of uncertain endpoints (results not shown here). The procedure did very well at estimating the fixed effects (including variable selection), regardless of the quality of the baseline hazard estimate. The random effects along with the level and shape of the hazard were also estimated very well (i.e., posterior distributions were tightly clustered around the true values) unless the sample size was small (less than 100) and the log-hazard deviated much from quadratic. For small samples, credible sets for random effects and the hazard function become much wider, but still reliably contained the true values, provided that somewhat reasonable priors were used.

4 Analysis of the Beam Testing Data

For the analysis of the beam testing data, $x_{i,1}, \ldots, x_{i,6}$ are indicators for whether or not (1 or 0) the application running during the i-th trial was cg, corr, hpl, int_add, vpic, or varied, respectively, defined in Section 1.1. If all $x_{i,j} = 0$ for $j = 1, \ldots, 6$ then the component was in an idle condition. The covariate $x_{i,7}$ is an indicator for beam aim (0 = Cell, 1 = Opteron), and $x_{i,8}$ equals the beam width (1-inch or 2-inch). Each fixed effect parameter in the vector $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_8)$ was treated as a separate group, i.e., each could be independently selected for inclusion in the model. This makes sense for β_7 and

β_8 since they are representing different factors. For $\beta_1, \ldots, \beta_6$ the rational was that some applications may have some additional effect relative to the baseline Operating System behavior (idle condition), whereas other applications may not. Thus, a unique indicator for variable selection (3) is attached to each β_j.

The values of $z_{i,1}, z_{i,2}, z_{i,3}$ provide indicators for which Triblade is being tested (Triblade 1, 3, or 4), and $\gamma_1, \gamma_2, \gamma_3$ are grouped together as one factor in the analysis. That is, either there is positive random effect variance λ_1 for the entire Triblade factor or not.

4.1 Prior Specifications

Values for $\{\dot{a}_j\}_{j=1}^8$, $\{\dot{b}_j\}_{j=1}^8$, $\dot{c}_1$, $\dot{d}_1$, $\{\dot{e}_m\}_{m=0}^2$, $\{\dot{f}_m\}_{m=0}^2$, $\dot{g}$, $\dot{h}_1$, $\dot{h}_2$, $\dot{i}_1$, $\dot{i}_2$, $\dot{j}_1$, and $\dot{j}_2$, need to be chosen to complete the prior specification. The values chosen for this application have been carefully specified based on previous studies and expert solicitation. Section 4.4 provides a sensitivity analysis of the effect of these choices on the analysis.

The prior on β_j implies that it is identically zero with probability $1-\dot{a}_j$. Here $\dot{a}_j = 0.5$ is specified for each j to allow each fixed effect covariate a 50% chance of being in the model *a priori*. Neutron-induced soft error experts at LANL suggested that if any of the fixed effect covariates had an effect, a 10× multiple to the baseline hazard would be a very large effect. Therefore, when $\dot{b}_j = 1$, it is *a priori* assumed that the 99-th percentile of the multiplicative effect to the hazard function is 10. For the random effect due to differences among Triblades, $\dot{c}_1$ was set to 0.5 to allow the Triblade effect a 50% chance of being in the model *a priori*. LANL experts suggested there may be about a 20-30% difference between two Triblades (if there was a non-negligible difference). Therefore, the median of the HC prior on ψ_1, the random effect variance when it is positive, was set to $\dot{d}_1 = 0.4$. This gives a median relative difference between two randomly drawn Triblades (i.e., the ratio, largest to smallest, of their respective multipliers to the hazard) of about 1.3 *a priori*, if there is a difference. The value for the largest exposure before a constant hazard of t^* from (6) was set to 2.752×10^9 neutrons/cm^2, as this is an upper bound on

the amount of exposure that any of the components received.

A previous study on the Cell processor indicated a projected mean fluence until neutron induced soft error of 2.25×10^8 neutrons/cm^2 for the two Cells in the Cell beam aim, which corresponds to a failure rate (if it was constant) of 4.4×10^{-9} cm^2/neutron. However, it was indicated that this projection may only be accurate to within an order of magnitude. Hence $\dot{e}_0 = 2.5$ and $\dot{f}_0 = 1$ are used so that with $t^* = 2.752 \times 10^9$ neutrons/cm^2, the baseline (corresponding to Cells running in the idle condition) hazard "mean" level ϕ_0 is geometrically centered around 4.4×10^{-9} and contained within $(0.34, 58.2) \times 10^{-9}$ with 99% probability, *a priori*.

Further, the values $\dot{e}_1 = \dot{e}_2 = 0$ and $\dot{f}_1 = \dot{f}_2 = 0.25$ are used, to encourage the inclusion of linear plus quadratic terms if they are needed. The median of the HC prior on τ is set to $\dot{g} = 0.3$, to encourage small values of τ so that departures from the quadratic trend for the log-hazard should be small unless the data suggest otherwise.

The prior used for κ (the proportional multiplier to get from number of bit flips to fluence for the Virtex-IIs) is centered around the sensitive area (the total area of the bits on the Virtex II) of 3.586×10^{-8} cm^2, scaled by $t^* = 3.725 \times 10^{10}$ to result in a mean for κ of 1336. Thus the values $\dot{h}_1 = 668$ and $\dot{h}_2 = 0.5$ are used to suggest that κ is in the interval (1219, 1459) with 99% probability *a priori*.

A fairly diffuse prior was used for μ_α, with $\dot{i}_1 = 3.35$ and $\dot{i}_2 = 10$. This implies that the proportion of neutrons remaining after passing through the Triblade is anywhere from $(0.04, 0.99)$ with 99% probability, *a priori*. The vague prior for μ_α is also justified in the sense that there is very little posterior sensitivity to this prior specification anyhow. The prior for σ_α^2 was set according to the assertion by LANL experts that the difference in reduction due to attenuation between different beam aims should be less than 10-20%. Therefore, we set $\dot{j}_1 = 1$ and $\dot{j}_2 = 200$ so that the standard deviation σ_α is anywhere between $(0.011, 0.136)$ with 95% probability, *a priori*.

4.2 Assessing the Adequacy of Model Assumptions

Before proceeding to the analysis, some of the key assumptions of the model presented in Section 2 are assessed. Four variables that could affect the hazard are included in this analysis, here referred to as factors since they take a small number of discrete values, (i) Application (cg, corr, hpl, int_add, varied, vpic, idle), (ii) Triblade (1, 3, 4), (iii) Beam Aim (Opteron, Cell), and (iv) Beam Width (1-inch, 2-inch). To assess the proportional hazards assumption for each factor, the data were split into subsets depending on the level of a particular factor. For a given factor the model was then fit with all of the other covariates included for each subset. Finally, the posterior median of the logarithm of the baseline hazard $log(h_0(y))$ corresponding to the model fit on each subset (one for each level of the given factor) of the data was plotted (not shown). The corresponding log-hazard in each plot showed no significant deviation from parallel lines, indicating that the proportional hazards assumption is plausible in this case.

The probability integral transform was also used in checking overall model fit. If (a_i^*, b_i^*), $i = 1, \ldots, n$, is an interval censored sample from the model in Section 2, then $(S_i(b_i^*), S_i(a_i^*))$ is an interval censored sample from a Uniform$(0, 1)$ distribution. The conditional distribution of a Uniform$(0, 1)$ random variable, given that it is in the interval $(S_i(b_i^*), S_i(a_i^*))$, is Uniform$(S_i(b_i^*), S_i(a_i^*))$. So for each i, a draw $U_i \sim$ Uniform$(S_i(b_i^*), S_i(a_i^*))$, $i = 1, \ldots, n$ from the posterior distribution is made. If the model is a reasonable approximation for these data, then the U_i, $i = 1, \ldots, n$, should be characteristic of an *iid* sample from a Uniform$(0, 1)$ distribution and $Z_i = \Phi^{-1}(U_i)$ (where Φ is the standard normal CDF) should be characteristic of an *iid* sample from a standard normal distribution. Figure 2(a) displays a normal Q-Q plot of the Z_i, while Figure 2(b) presents the Z_i plotted against the order of the experimental trials. The Q-Q plot shows that the points fall very close to the $y = x$ line, indicating good agreement between the model and the data. Thus, there appears to be no systematic bias in the es-

Figure 2: Probability Integral Transform Residual Diagnostic Plots.

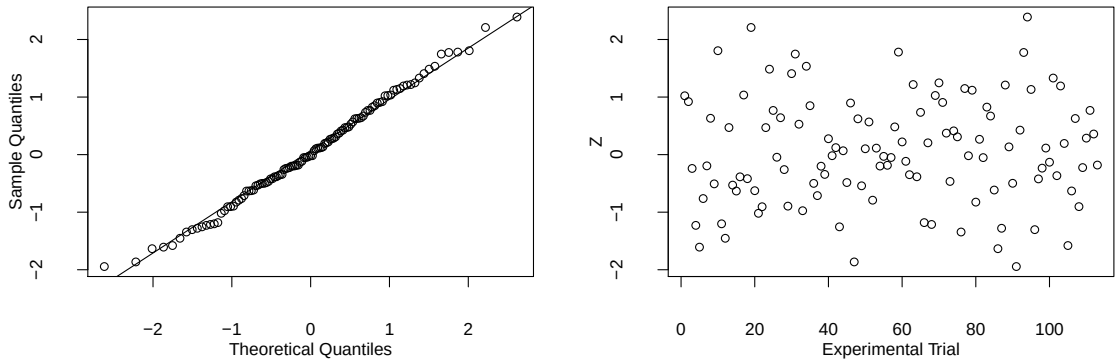

(a) Normal Q-Q plot of Z residuals

(b) Z residuals across experimental trial order

timate of the distribution depending on the value of exposure, leading to confidence that the model is sufficiently flexible for this data set. The plot on the right shows random scatter across the order of the experimental trials, indicating no trend in the Z_i deviates with the sequential order in which the trials were performed.

4.3 Summary of the Fitted Model

It appeared (based on time series plots) that the MCMC Algorithm described in the Supplemental Material was in its steady state for this application somewhere before 10,000 iterations, so 50,000 posterior samples were collected, starting with iteration 10,001. The total 60,000 iterations took about 12 hours on a MacBook Pro with a 3.06 GHz Intel Core 2 Duo processor. The results below reflect cosmic-ray-neutron-induced errors (failures and SDCs), and do not include those due to other causes. Further, all results are specific to the experimental conditions, e.g. angle of incidence of the beam, and are based on the experimental data and the model used for the data and do not neccessarily reflect failures or SDCs observed in the Roadrunner platform.

Figure 3 displays 50 randomly selected posterior draws of the baseline hazard function for the full model, along with the posterior median and 95% credible bands. The median hazard appears mostly flat like an exponential failure distribution, and it is clear that a

constant curve fits well within the 95% bands, thus there is no evidence of a departure from a constant hazard. However, the 95% bands are fairly wide, and a log-linear or log-quadratic trend could easily fit within these bands as well, and some individual posterior draws of the hazard do exhibit some "bath-tub" curve behavior. Thus, it is certainly possible that the hazard may increase at larger levels of exposure. With additional data with even higher exposure counts, we speculate that there might be more evidence to suggest that the hazard rate eventually increases, but this is not much evidence to suggest that this is the case based on these data.

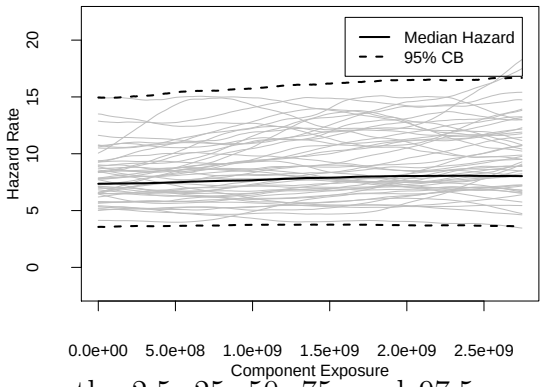

Figure 3: Baseline Hazard Function. Posterior median hazard (solid) along with 95% Credible Bands (dashed) and 50 randomly selected posterior draws (grey).

The posterior distribution results for the fixed effects and the random effects are summarized in Table 1. It is preferable to use the median (as opposed to the mean) as a summary measure of the "center" of the posterior distribution in this analysis, since the distribution of some of the parameters (particularly those that are related in some way to a HC prior) have fairly heavy right tails. Thus, for consistency, the 2.5, 25, 50, 75, and 97.5 percentiles are presented, respectively, along with the probability that the effect is nonzero where appropriate, to summarize the posterior distribution of fixed effects and random effects parameters.

The application with the potentially largest effect is hpl, but it only has posterior probability of a non-zero effect (i.e., different from the idle condition) of 0.417. The posterior median of the hazard multiplier due to the hpl application is 1.000 with a 95% credible interval (CI) of $(1.000, 2.545)$, indicating at most a modest increase in the rate of errors for this application. The application in this study that is most similar to an application that would get used in the field is vpic. The posterior probability that vpic has a different error rate than idle is 0.202, with a 97.5 percentile for the multiplicative

Table 1: Posterior Summary of Fixed Effects and Random Effects Parameters: $\exp\{\beta\}$ is the multiplicative effect to the hazard for the given effect, while $\exp|\gamma - \gamma'|$ is the relative difference in hazard rate between two randomly selected Triblades (larger to smaller).

Fixed Effect	$\Pr(\beta \neq 0)$	Posterior Summary of $\exp\{\beta\}$				
		2.5%	25%	50%	75%	97.5%
cg (β_1)	0.195	0.654	1.000	1.000	1.000	1.457
corr (β_2)	0.158	0.723	1.000	1.000	1.000	1.247
hpl (β_3)	0.417	1.000	1.000	1.000	1.494	2.545
int_add (β_4)	0.594	0.237	0.452	0.720	1.000	1.000
varied (β_5)	0.240	0.581	1.000	1.000	1.000	1.681
vpic (β_6)	0.202	0.600	1.000	1.000	1.000	1.356
Beam Aim (β_7)	1.000	2.749	4.574	5.884	7.525	11.753
Beam Diam (β_8)	0.198	0.675	1.000	1.000	1.000	1.468

Random Effect	$\Pr(\lambda \neq 0)$	Posterior Summary of $\exp\|\gamma - \gamma'\|$				
		2.5%	25%	50%	75%	97.5%
Triblade (λ_1)	0.897	1.000	1.098	1.357	1.946	5.050

effect of 1.356. Overall these results indicate that there is not very much evidence that the hazard rate differs for the various applications versus an idle condition. While this result may seem surprising, it is not unreasonable, as the processor is not literally idle when it is not running an application. In particular it is still executing instructions in response to the operating system. Of course, with additional time on test or with other applications, different results might be found.

The coefficient for Beam Aim (β_7) has a posterior probability of 1.00 that $\beta_7 \neq 0$, indicating a definite difference between Cell beam aim and Opteron beam aim. The posterior distribution indicates a median multiplier to the hazard rate of 5.884 with a 95% CI of (2.749,11.753), meaning there is roughly six times more risk of an error when the beam is aimed at Opterons as opposed to Cells. This result is not the same as saying that the Opteron has a hazard rate that is roughly six times the Cell hazard rate. Other hardware along the trajectory of the beam when it is aimed at an Opteron or hardware affected by scatter may be responsible for errors that resulted when the beam was aimed at an Opteron. The same holds for the Cell beam aim so that the root cause of errors in

either beam aim cannot be conclusively determined.

There is partial confounding between the Triblade under test and the beam width used for the testing. Specifically, the 2-inch beam was used for testing Triblade 3 , the 1-inch beam was used for testing Triblade 4 and both were used for testing Triblade 1. In a situation such as this, it can be hard for the model to determine whether it is the Triblade under test or the beam width that has a more influential effect on the hazard rate. Nonetheless, the posterior probability that either the Triblade under test or the beam width (or both) is an important predictor of the hazard rate is 0.93, and the results below suggest that it is likely the Triblade under test that has the larger impact on the hazard rate.

The posterior probability that coefficient β_8 (representing the beam width effect) differs from zero is 0.198. That is, there is not much evidence to suggest that beam width alone has an effect on fluence until error distribution. The 97.5 percentile of the multiplicative effect is 1.468, indicating that even if beam width does have an effect it is not very large. This means that the error rate for a 1 inch width beam is not that different from a 2 inch width beam, suggesting that the components most crucial to failure were likely covered by the 1 inch beam width.

The posterior probability that $\lambda_1 \neq 0$ is 0.897, which indicates that there is likely a nonzero random effect due to Triblade. The posterior median of λ_1 is 0.381 with a 95% CI of (0.000, 1.294) (not shown in Table 1). This leads to a median relative difference (largest to smallest) between two Triblades of about 1.357 with a 95% CI of (1.000, 5.050). The effect due to Triblade appears to be the second largest effect in this study after that of Beam Aim.

4.4 Prior Sensitivity

Variable selection in our model (i.e., proportion of nonzero draws in the posterior) can be sensitive to the prior placed on the βs for fixed effects and λs for random effects.

For example, a more diffuse prior for the positive part of βs and λs can lead to a lower probability of inclusion. However, general summaries, such as 95% CIs for example, for the βs and γs do not seem to be that sensitive to the prior placed on these parameters.

The shape of the baseline hazard h_0 for individual draws, and to a lesser extent the "mean" level (ϕ_0) of h_0, are somewhat sensitive to their prior specifications. This is because this curve has a lot of flexibility in the model of Section 2.1. In addition, there is diminishing information about this curve at higher exposures. Summaries such as the posterior median h_0 and 95% confidence bands, for example, have some sensitivity to prior assumptions, but less so than the extremes of the distribution. With more diffuse priors on the constant, linear and quadratic trends and g_0 standard deviation, ϕ_0, ϕ_1, ϕ_2, and τ, respectively (i.e., with $\dot{e}_0 = \dot{e}_1 = \dot{e}_2 = 0$, $\dot{f}_0 = 100$, $\dot{f}_1 = \dot{f}_2 = 10$, and $\dot{g} = 1$, as compared to the priors used in the analysis of $\dot{e}_0 = 5, \dot{e}_1 = \dot{e}_2 = 0$, $\dot{f}_0 = 1$, $\dot{f}_1 = \dot{f}_2 = 0.25$, and $\dot{g} = 0.3$), the median hazard and 95% confidence bands only had relative changes of about 25%.

The posterior distributions of the fixed effect parameters in the model, i.e., the βs, do not appear to be sensitive at all to changes in the hazard priors. The random effect parameters do have some sensitivity to the changes in hazard priors, due to the confounding that is inherent between the γs (random effect coefficients) and the ϕs (the quadratic trend coefficients of the log-hazard). For example, suppose a fairly diffuse prior is assumed for the random effect standard deviation λ, and a tight prior which is not consistent with the data is assumed for the hazard. Then in order to pull the parameter values away from the prior toward the data, λ will be made large and the γs will be made either all large and positive or all large and negative as needed to account for the poor prior on the hazard function. In this case the posterior medians of γ_1, γ_2, and γ_3, were 0.17, 0.00, and -0.43, respectively, with a posterior median for the sum of the three γs of -0.201, which indicates some possible biasing effect of the hazard like that mentioned above, but also no substantial biasing effect.

The beam attenuation model parameters κ and μ_α have very little sensitivity to prior specification within any reasonable bounds. This is due to the large number of Virtex-II observations, each of which provides direct information about these parameters. The posterior distributions of the α_r do have some sensitivity to the prior on σ_α^2, however. For example if set very tightly near zero, then all α_r are tightly centered around μ_α, but if set very diffuse, then the α_r become much more diffuse. Thus, it was important to get expert solicitation involved in setting this prior. Also, while the α_r show some sensitivity to the prior on σ_α^2, the end results are not altered substantially, and the conclusions in this section remain unchanged for any reasonable prior, which gives us further confidence in their validity.

4.5 Projection of Failures and SDCs to Clock Time

Until now all errors (SDC or failure) have been treated as one collective class of errors. Now SDCs are considered separately from all errors that are not SDCs, i.e., detectable errors such as crashes, are denoted as *failures*. To this end, it is necessary to model the probability that a given error is an SDC as opposed to a failure.

It is assumed that the probability that a given error is an SDC, Pr(SDC), is constant over the levels of neutron exposure in this study, independent among trials, and not affected by the covariates with the exception of Beam Aim. There is some speculation that the other covariates could affect Pr(SDC), e.g., some applications could produce more SDCs than others. However, an effect on exposure until any error due to any of the applications was not apparent fromn the data, and the above assumptions were also said to be reasonable by LANL experts. Further, with so few SDCs (four total in the entire experiment), it would be very difficult to capture these relationships with any accuracy.

With the above assumptions, Pr(SDC) is estimated independently from the parameters of the model in Section 2 based on the binomial data from the errors that occurred in the experimental trials. That is, let N^o be the number of errors observed when the

beam was aimed at an Opteron that could possibly have been a SDC (as explained in the next paragraph) and let N^o_{SDC} be the number of those errors that were SDCs. Denote the probability that a given error from an Opteron is an SDC as $\Pr^o(\text{SDC})$. Finally, let N^c, N^c_{SDC}, and, $\Pr^c(\text{SDC})$ be similarly defined for the Cells. Based on expert judgment, a Beta$(2, 20)$ prior on both $\Pr^o(\text{SDC})$ and $\Pr^c(\text{SDC})$ is assumed, which implies that the proportion of all errors that are SDCs is anywhere from 0.005 to 0.304, with 99% probability, *a priori*. The posterior distributions of $\Pr^o(\text{SDC})$ and $\Pr^c(\text{SDC})$ are then Beta$(N^o_{SDC} + 2, N^o + 20)$ and Beta$(N^c_{SDC} + 2, N^c + 20)$, respectively.

In this study $N^o = 11$, $N^o_{SDC} = 2$, $N^c = 76$, and $N^cSDC = 2$. The total of $N^o + N^c = 87$ is smaller than the total number of trials $n = 113$. This is because not every trial ended in an error (some were right censored) and not every error could have been an SDC (some errors occurred in the idle state, and an SDC could not have been detected unless the trial had a computational application running). This results in a posterior median for $\Pr^o(\text{SDC})$ of 0.114, with a 95% CI of (0.035, 0.250), and a posterior median for $\Pr^c(\text{SDC})$ of 0.038, with a 95% CI of (0.011, 0.088).

Based on the estimated model, predictive distributions for the clock time until a failure and the clock time until SDC that might be observed in practice for each of the two beam aims (hardware in the Opteron beampath or Cell beampath, respectively) are presented in Figure 5. These distributions are for the application "vpic" as it is the most realistic example (of the applications used in this study) of a code that would be run in the field, although the applications in this study did not demonstrate much difference from one another anyhow.

These results have been calculated for the neutron flux in Los Alamos, NM, which is estimated to be approximately five times that at sea level (`http://seutest.com`), and are based on fluences for neutrons with energies above 10 MeV. The 10 MeV cutoff is a value suggested by (JEDEC Solid State Technology Association 2001) as neutrons with energy much below 10 MeV are not thought to be able to cause bit flips, but the lowest

energy (ε) required to cause a soft error is component dependent. In any case, the exact cutoff used here would not affect projections provided the neutron spectrum at the ICE House was the same as that for cosmic-ray-induced neutrons in Los Alamos above the neutron energy corresponding to the unknown ε. In Figure 4 it can be seen that the spectrums are not identical, particularly near the low end and for very high neutron energies (Wender 2003). However, they are very similar near and above 10 MeV (i.e., near where ε is believed to be) until very high values of neutron energies are reached, where there are so many fewer neutrons that it can be assumed that the difference has a negligible effect on projections.

Figure 4: Comparison of LANSCE neutron spectrum with measured cosmic-ray induced neutron spectrum in Los Alamos.

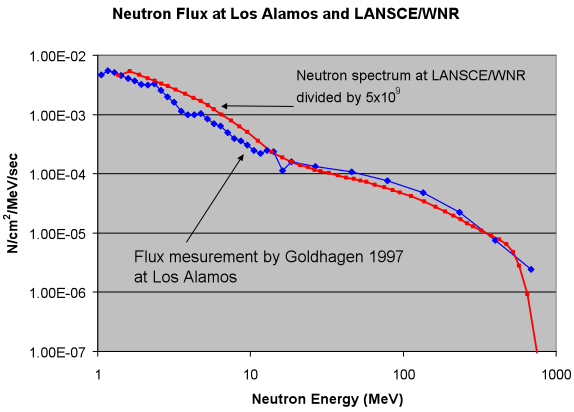

There is also variability in the neutron flux over time in Los Alamos. To account for this, it is assumed, that the ambient neutron flux in Los Alamos, NM is normal with mean 0.019 neutrons/cm^2/sec and standard deviation 4.4×10^{-4} neutrons/cm^2/sec. This assumption is the same as that used in a previous study (Michalak et al. 2005) only with the mean used in that work (0.025 neutrons/cm^2/sec) corrected to the value obtained according to `http://seutest.com`. Thus, the fluence over time is modeled as a Brownian Motion (BM) path with drift equal to 0.019 neutrons/cm^2/sec and scale factor equal to 4.4×10^{-4} neutrons/cm^2/sec.

The predictive distributions for exposure until failure and SDC for Cell and Opteron beaim aims in Figure 5 are produced by generating 10,000 draws (then smoothing them) from the following Monte Carlo simulation procedure.

Algorithm 1.

1. *Randomly draw a set of parameters from the posterior distribution.*

2. *Randomly draw a Triblade effect* $\gamma_k \sim N(0, \tau^2)$ *given the sampled value of* τ *in 1.*

3. *Randomly generate an exposure until error* E *according to the hazard model in (2) with the parameters given in step 1 and application = "vpic", beam width = 2 inches, and beam aim specific to the case at hand. If simulating SDCs use a baseline hazard function of* $h_0 \times \Pr^o(SDC)$ *for Opteron Beam aim or* $h_0 \times \Pr^c(SDC)$ *for Cell Beam aim (where* $\Pr^o(SDC)$ *or* $\Pr^c(SDC)$ *is obtained from step 1) as the baseline hazard function. If simulating failures use a baseline hazard of* $h_0 \times [1 - \Pr^o(SDC)]$ *(or* $h_0 \times [1 - \Pr^c(SDC)]$*). This is equivalent to the assumptions made above for the occurrence of SDCs and failures.*

4. *Randomly generate a BM path with drift equal to 0.019 neutrons/cm*^{2}*/sec and scale factor equal to* 4.4×10^{-4} *neutrons/cm*^{2}*/sec, until the time* T *when it is first greater than or equal to the exposure* E *produced in 3. Return* T*.*

The predictive distributions in Figure 5 appear somewhat similar to the exponential distribution in shape, but they are actually quite a bit more heavy tailed. This is because the hazard function uncertainty is included in these distributions, so they are actually more similar to an exponential distribution with a random rate (they would be exactly that if the hazard was assumed constant in our model).

The caption of Figure 5 provides the approximate posterior median of the mean time to failure (MTTF) or mean time to SDC (MTTS), along with 95% CIs for the MTTF (or MTTS) for each case. The posterior median and CIs for the MTTFs are produced by calculating the mean of the time until failure (or SDC for MTTS) distribution (by numerical integration) for each of 10,000 draws from the posterior distribution parameters as in steps 1 and 2 of Algorithm 1.

Further, the modeling results were used to estimate the predictive distributions of time until failure and time until SDC for a single Triblade and for the 180 compute Triblades in a single connected unit (CU) (Roadrunner is composed of 17 CUs); see Figure 6. These distributions only incorporate hardware in the Cell and the Opteron beampaths, so do not include *all* hardware in a Triblade. Specifically, independence of the errors occurring at the four different beam aims, whether Opteron (one for each of

Figure 5: (a) Predictive Distribution of the Time until a Failure for all hardware in the Cell beampath. The posterior median for the mean time to failure (MTTF) is 2.42×10^5 days with a 95% CI of $(0.79, 4.52) \times 10^5$ days. (b) Predictive Distribution of the Time until a Failure for all hardware in the Opteron beampath. The posterior median MTTF is 4.43×10^4 days with a 95% CI of $(1.42, 13.62) \times 10^4$ days. (c) Predictive Distribution of the Time until a SDC for all hardware in the Cell beampath. The posterior median for the mean time to SDC (MTTS) is 5.83×10^6 days with a 95% CI of $(1.60, 19.48) \times 10^6$ days. (d) Predictive Distribution of the Time until a SDC for all hardware in the Opteron beampath. The posterior median MTTS is 3.49×10^5 days with a 95% CI of $(0.92, 14.10) \times 10^5$ days.

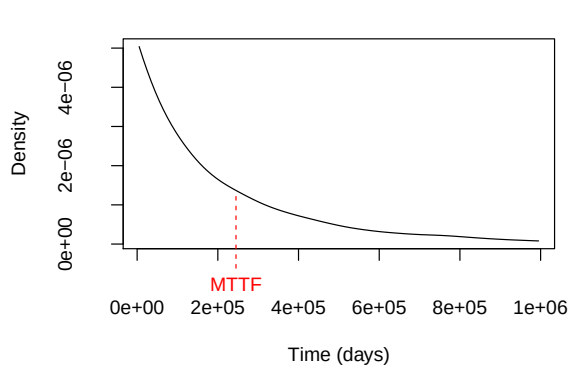

(a) Time to Failure for Cell Beam Aims

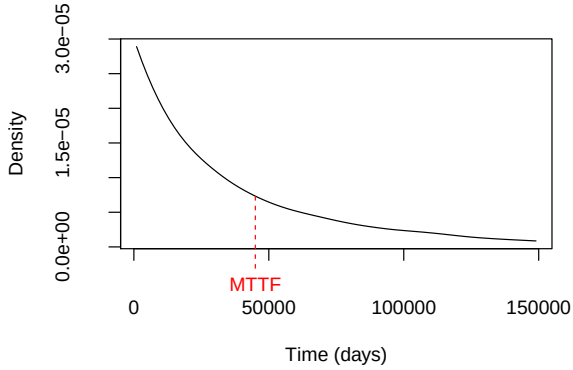

(b) Time to Failure for Opteron Beam Aims

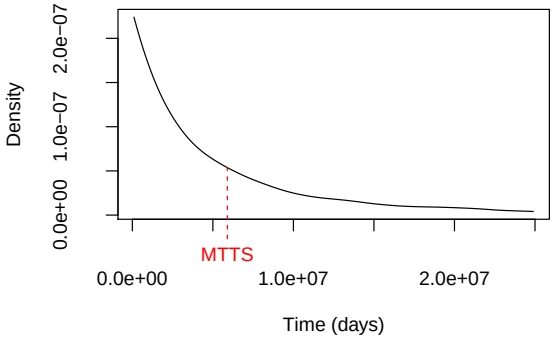

(c) Time to SDC for Cell Beam Aims

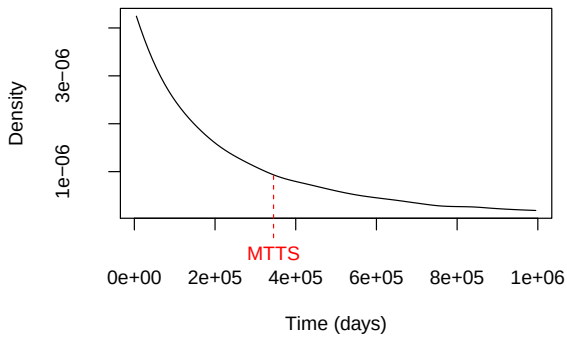

(d) Time to SDC for Opteron Beam Aims

the two Opterons in the LS21 blade) or Cell (one for each of the two Cells in the QS22), is assumed. Therefore, the hazard functions for these four distinct aims are added together to obtain an approximate hazard for an entire Triblade. Independence among Triblades is assumed to get an overall hazard function for a CU in a similar manner.

The distributions for exposure until failure and SDC in Figure 6 were generated in the same fashion as in Algorithm 1, except that the hazard for simulating E in step 3 is obtained by adding up the respective hazards for the four different beam aims to get

Figure 6: (a) Predictive Distribution of time until a Failure for one Triblade. The posterior median MTTF is 1.87×10^4 days with a 95% CI of $(0.59, 5.17) \times 10^4$ days. (b) Predictive Distribution of of time until a Failure for one Roadrunner CU. The posterior median MTTF is 92 days with a 95% CI of $(37, 182)$ days. (c) Predictive Distribution of time until a SDC for one Triblade. The posterior median MTTS is 1.62×10^5 days with a 95% CI of $(0.42, 4.58) \times 10^5$ days. (d) Predictive Distribution of of time until a SDC for one Roadrunner CU. The posterior median MTTS is 796 days with a 95% CI of $(258, 2448)$ days.

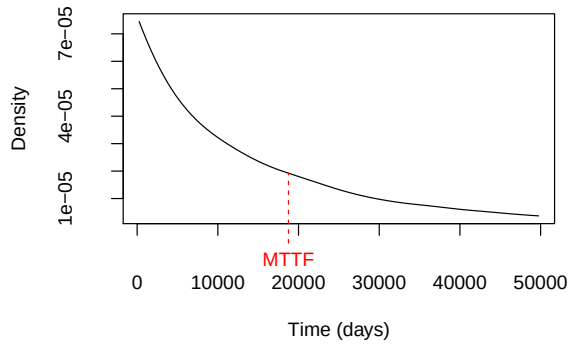

(a) Time to Failure for a single Triblade

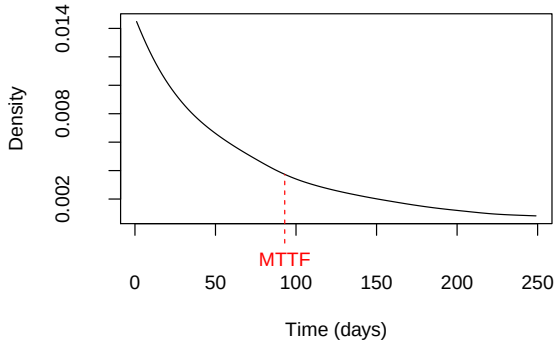

(b) Time to Failure for a single Connected Unit

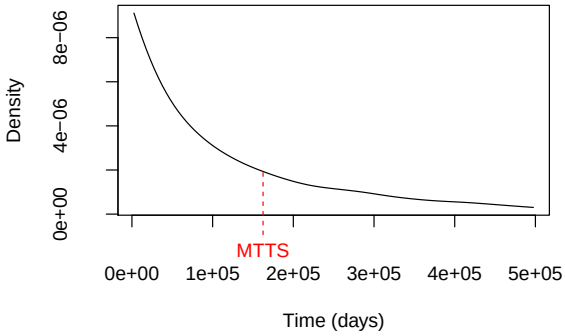

(c) Time to SDC for a single Triblade

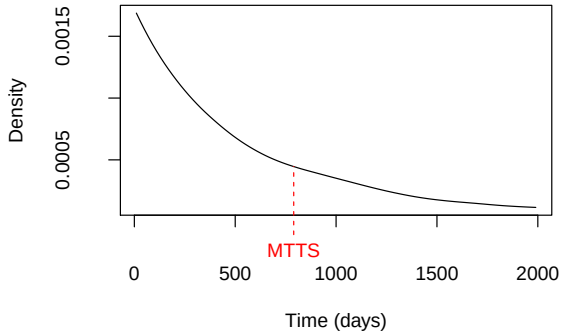

(d) Time to SDC for a single Connected Unit

the hazard for a single Triblade. The hazard to generate E for a CU (180 Triblades) is obtained by drawing parameters as in step 1 of Algorithm 1, but then each of 180 Triblade hazards are calculated with a different random Triblade effect from step 2 (all 180 of these are drawn from the value of τ drawn in step 1). These 180 Triblade hazards are then added together to get the hazard to generate a failure time for a CU (and analogously for generating SDCs).

As mentioned earlier, these distributions reflect the approximate neutron flux in Los Alamos, NM, and are subject to errors due to beam fluence measurement inacurracy, as

well as any potential lack of model fit. The Triblade and CU time to failure and time to SDC information is extrapolated from the neutron beam testing and does not represent observed cosmic-ray-neutron-induced failures or SDCs in the Roadrunner supercomputer (such data are unavailable).

Even with the large uncertainty, these results indicate that a failure on a CU is very likely to occur within six months, and we would expect failures to occur roughly once every three months. An SDC would be expected to occur on a CU every 26 months, and it is very likely for an SDC to occur on a CU within four years. On the other hand, SDCs could be occurring on a CU as often as once every eight months according to these results Of course, these results are specific to the vpic code and and all other conditions during the testing, e.g., angle of incidence of the neutron beam, fluence measurement accuracy etc., as well as to all of the modeling assumptions.

5 Conclusions

Accelerated neutron beam testing was performed on components of Roadrunner to assess the importance of certain variables on the effects of cosmic-ray-induced neutrons on the system. A novel statistical approach for the analysis of interval censored survival data with mixed effects and uncertainty on the interval endpoints was developed for the analysis of these data. A key result is that the system has a substantially elevated hazard ($\sim 6\times$) when the beam was aimed at Opterons as opposed to Cells. Estimates based on this accelerated testing of the distributions of failures and of SDCs in actual clock times for the neutron flux in Los Alamos were provided at the component level (Cell beam aims or Opteron beam aims), Triblade, and a Roadrunner CU. Results indicate an estimated mean time to failure for a Roadrunner CU of about 92 days, and a mean time to SDC of about 796 days, with roughly an order of magnitude of uncertainty around these estimates.

Acknowledgments

Many people contributed to the success of the Triblade beam-testing experiments. The authors thank Bill Rust, David Modl, Andrea Manuzzato, Sean Blanchard, Joe Abeyta, Chuck Alexander, Ben Bergen, Ann Borrett, Henry Brandt, James Campa, Randy Cardon, Nathan DeBardeleben, Tom Fairbanks, Parks Fields, Alan Gibson, Gary Grider, Josip Loncaric, Pablo Lujan, Alex Malin, Fred Marshall, Andrew Montoya, John Morrison, Andrew Shewmaker, Manuel Vigil, Bob Villa, Cornell Wright, and the reviewers of the IEEE TDMR paper documenting some more of these results, and apologize for any inadvertent omissions from this list.

References

Abrahamowicz, M., Ciampi, A. & Ramsay, J. (1992), 'Nonparametric density estimation for censored survival data: Regression-spline approach', *The Canadian Journal of Statistics* **20**, 171–185.

Ando, H., Kan, R., Tosaka, Y., Takahisa, K. & Hatanaka, K. (2008), 'Validation of hardware error recovery mechanisms for the sparc64 v microprocessor', *Proceedings of the 2008 IEEE International Conference on Dependable Systems and Networks* pp. 62–69.

Baumann, R. (2005), 'Radiation-induced soft errors in advanced semiconductor technologies', *IEEE Transactions on Dev. and Mat. Rel.* **5**, 305–316.

Constantinescu, C. (2005), 'Neutron ser characterization of microprocessors', *Proceedings of the 2005 International Conference on Dependable Systems and Networks* pp. 754–759.

Cox, D. (1972), 'Regression models and life-tables', *Journal of the Royal Statistical Society. Series B (Methodological)* **34**, 187–220.

Fan, J. & Li, R. (2001), 'Variable selection via nonconcave penalized likelihood and its oracle properties', *Journal of the American Statistical Association* **96**, 1348–1360.

Fan, J. & Li, R. (2002), 'Variable selection for cox's proportional hazards model and frailty', *Annals of Statistics* **30**, 74–99.

Faraggi, D. & Simon, R. (1998), 'Bayesian variable selection method for censored survival data', *Biometrics* **54**, 1475–1485.

Finkelstein, D. (1986), 'A proportional hazards model for interval-censored failure time data', *Biometrics* **42**, 845–854.

Gelfand, A. (1990), 'Sampling-based approaches to calculating marginal densities', *Journal of the American Statistical Association* **85**(410), 398–409.

Gelman, A. (2006), 'Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper)', *Bayesian Analysis* **1**, 515–534.

Gelman, A., Carlin, J., Stern, H. & Rubin, D. (2004), *Bayesian Data Analysis*, 2nd edn, Chapman & Hall, London.

George, E. & McCulloch, R. (1993), 'Variable selection via gibbs sampling', *Journal of the American Statistical Association* **88**, 881–889.

George, E. & McCulloch, R. (1997), 'Approaches for Bayesian variable selection', *Statistica Sinica* **7**, 339–373.

Goetghebeur, E. & Ryan, L. (2000), 'Semiparametric regression analysis of interval-censored data', *Biometrics* **56**, 1139–1144.

Goggins, W., Finkelstein, D., Schoenfeld, D. & Zaslavsky, A. (1998), 'A markov chain monte carlo em algorithm for analyzing interval-censored data under the cox proportional hazards model', *Biometrics* **54**, 1498–1507.

Hong, T., Michalak, S., Graves, T., Ackaret, J. & Rao, S. (2009), 'Neutron beam irradiation study of workload dependence of ser in a microprocessor', *2009 SELSE Proceedings* .

Ibrahim, J., Chen, M. & MacEachern, S. (1999), 'Bayesian variable selection for proportional hazards models'.

JEDEC Solid State Technology Association (2001), 'JEDEC standard JESD89: Measurement and reporting of Alpha particles and terrestrial cosmic-ray-induced soft errors in semiconductor devices', *Scientific Computing* . [Online].
URL: *http://www4.tsl.uu.se/ bumpen/jedec.pdf*

Kinney, S. & Dunson, D. (2007), 'Fixed and random effects selection in linear and logistic models', *Biometrics* **63**, 690–698.

Koch, K. (2008), 'Roadrunner platform overview'. [Online].
URL: *http://www.lanl.gov/orgs/hpc/roadrunner/pdfs/Koch%20-%20Roadrunner%20Overview/RR%20Seminar%20-%20System%20Overview.pdf*

Kooperberg, C. & Clarkson, D. (1996), 'Hazard regression with interval-censored data', *Biometrics* **53**, 1485–1494.

Lee, K., Chakraborty, S. & Sun, J. (2011), 'Bayesian variable selection in semiparametric proportional hazards model for high dimensional survival data'. Unpublished Manuscript.

Meuer, H. (2008), '31st top500 list topped by first-ever petaflop/s supercomputer', *Scientific Computing* . [Online].
URL: *http://www.scientificcomputing.com/31st-TOP500-List-Topped-by-First-ever-Petaflops-Supercomputer.aspx*

Michalak, S., DuBois, A., Storlie, C., Quinn, H., Rust, W., DuBois, D., Modl, D., Manuzzato, A. & Blanchard, S. (2011), 'Neutron beam testing of high performance computing hardware', *2011 NSREC Radiation Effects Data Workshop Proceedings, to appear* .

Michalak, S., DuBois, A., Storlie, C., Rust, W., DuBois, D., Modl, D., Quinn, H., Manuzzato, A. & Blanchard, S. (2011), 'Neutron beam testing of triblades', *2011 SELSE Proceedings* .

Michalak, S., Harris, K., Hengartner, N., Takala, B. & Wender, S. (2005), 'Predicting the number of fatal soft errors in los alamos national laboratorys asc q supercomputer', *IEEE Trans. Dev. and Mat. Rel.* **5**.

O'Hara, R. & Sillanpää, M. (2009), 'A review of bayesian variable selection methods: What, how and which', *Bayesian Analysis* **4**, 85–118.

Raftery, A., Madigan, D. & Volinsky, C. (1995), Accounting for model uncertainty in survival analysis improves predictive performance, *in* 'In Bayesian Statistics 5', University Press, pp. 323–349.

Rao, S., Hong, T., Sanda, P., Ackaret, J., Barrera, A., Yanez, J., Mitra, S., Kellington, J. & McBeth, R. (2008), 'Examining workload dependence of soft error rates', *2008 SELSE Proceedings* . [Online].
URL: *http://www.selse.org*

Reich, B., Storlie, C. & Bondell, H. (2009), 'Variable selection in Bayesian smoothing spline ANOVA models: Application to deterministic computer codes', *Technometrics* **51**, 110–120.

Roberts, G., Gelman, A. & Gilks, W. (1995), 'Weak convergence and optimal scaling of random walk metropolis algorithms', *The Annals of Probability* **7**, 110–120.

Sanda, P., Kellington, J., Kudva, P., Kalla, R., McBeth, R., Ackaret, J., Lockwood, R., Schumann, J. & Jones, C. (2008), 'Soft-error resilience of the ibm power6 processor', *IBM Journal of Research and Development* **52**, 275–284.

So, Y., Johnston, G. & Kim, S. (2010), Analyzing interval-censored survival data with sas software, *in* 'Proceedings of the SAS Global Forum 2010 Conference', SAS Institute Inc.

Stein, M. (1999), *Interpolation of Spatial Data*, Springer-Verlag, New York, NY.

Takala, B. (2006), 'The ice house: neutron testing leads to more-reliable electronics, los alamos science'. [Online].
URL: *http://library.lanl.gov/cgi-bin/getfile?30-12.pdf*

Tibshirani, R. (1997), 'The lasso method for variable selection in the cox model', *Statistics in Medicine* **16**, 385–395.

Volinsky, C., Madigan, D., Raftery, A. & Kronmal, R. (1997), 'Bayesian model averaging in proportional hazard models: Assessing the risk of a stroke', *Applied Statistics* pp. 433–448.

Wahba, G. (1990), *Spline Models for Observational Data*, CBMS-NSF Regional Conference Series in Applied Mathematics.

Wender, S. (2003), Neutron single event effects testing at lansce, *in* 'IEEE International Reliability Physics Symposium'.

Xilinx (2007), 'Virtex-ii platform fpgas complete data sheet'. [Online].
URL: *http://www.xilinx.com/support/documentation/data_sheets/ds031.pdf*

Zhang, H. & Lu, W. (2007), 'Adaptive lasso for cox's proportional hazards model', *Biometrika* **94**, 691–703.

Ziegler, J. (1996), 'Terrestrial cosmic rays', *IBM Journal of Research and Development* **40**, 19–40.

Ziegler, J. & Lanford, W. (1981), 'The effect of sea-level cosmic rays on electronic devices', *Journal of Applied Physics* **52**, 4305–4312.

A Supplementary Material: MCMC Algorithm

The description below is written for the most general case of the methodology. All parameter values are initialized at the mode of the prior distribution, then the following steps are repeated until an adequate number of posterior samples have been obtained for convergence and inference.

MCMC Algorithm.

1. *For each fixed effects factor $l = 1, \ldots, g_f$, update $\boldsymbol{\beta}_l = \{\beta_j : j \in J_l\}$ simultaneously with Metropolis Hastings (MH). The proposal $\boldsymbol{\beta}_l^*$ given the current value $\boldsymbol{\beta}_l$ is generated for each $j \in J_l$ according to*

$$\beta_j^* = \begin{cases} \rho_{f,l}^* Z_{f,j}^* & \text{if } \boldsymbol{\beta}_l = \mathbf{0} \\ \varrho_{f,l}^* T_{f,j}^* & \text{otherwise,} \end{cases} \tag{A1}$$

where $\rho_{f,l}^* \sim$ Bernoulli$(\pi_{0,1})$, $\varrho_{f,l}^* \sim$ Bernoulli$(1-\pi_{1,0})$, $Z_{f,j}^* \overset{ind}{\sim} N(0, \dot{b}_j^2)$, $T_{f,j}^* \overset{ind}{\sim} \beta_j + \sigma_{f,j} t(\nu_{f,j})$, $t(\nu)$ is the t-distribution with ν degrees of freedom, and $\pi_{0,1}$, $\pi_{1,0}$, $\sigma_{f,j}$, and $\nu_{f,j}$, $j = 1, \ldots, p$ are control parameters that are adjusted to obtain good mixing. Here the following values are used: $\pi_{0,1} = 1$, $\pi_{1,0} = 0.25$, all $\nu_{f,j} = 10$ and the $\sigma_{f,j}$ range from 0.15 to 0.50 to encourage ∼30% acceptance of $\boldsymbol{\beta}_l^*$ when $\boldsymbol{\beta}_l \neq \mathbf{0}$ and $\varrho_{f,l}^* = 1$ in accordance with the recommendations in Roberts, Gelman & Gilks (1995) and Gelman, Carlin, Stern & Rubin (2004).

For each random effects factor $l = 1, \ldots, g_r$, update $\boldsymbol{\gamma}_l = \{\gamma_k : k \in K_l\}$ and $\boldsymbol{\lambda}_l = \{\lambda_k : k \in K_l\}$ simultaneously with MH. The proposals $\boldsymbol{\lambda}_l^*$, $\boldsymbol{\gamma}_l^*$ given the current values $\boldsymbol{\lambda}_l$, $\boldsymbol{\gamma}_l$ are generated for each $k \in K_l$ according to

$$\begin{aligned} \lambda_k^* &= \begin{cases} \rho_{r,l}^* C_l^* & \text{if } \boldsymbol{\lambda}_l = \mathbf{0} \\ \varrho_{r,l}^* U_l^* & \text{otherwise,} \end{cases} \\ \boldsymbol{\gamma}_k^* &= \begin{cases} Z_{r,k}^* I_{\{\lambda_k^* > 0\}} & \text{if } \boldsymbol{\lambda}_l = \mathbf{0} \\ T_{r,k}^* I_{\{\lambda_k^* > 0\}} & \text{otherwise,} \end{cases} \end{aligned} \tag{A2}$$

where $\rho_{r,l}^* \sim$ Bernoulli$(\pi_{0,1})$, $\varrho_{r,l}^* \sim$ Bernoulli$(1-\pi_{1,0})$, $C_l^* \overset{ind}{\sim} HC(\dot{d}_l)$, $U_l^* \overset{ind}{\sim} \exp\{\log \lambda_k + \varsigma_l t(\iota_l)\}$, where t is the student's t distribution, $Z_{r,k}^* \overset{ind}{\sim} N(0, \lambda_k^2)$, $T_{r,j}^* \overset{ind}{\sim} \gamma_k + \sigma_{r,k} t(\nu_{r,k})$, and $\pi_{0,1}$, $\pi_{1,0}$, ι_l, ς_l, $l = 1, \ldots, g_r$, $\sigma_{r,k}$ and $\nu_{r,k}$, $k = 1, \ldots, q$ are control parameters. The following parameter values are used: $\pi_{0,1} = 1$, $\pi_{1,0} = 0.25$, $\iota_l = 10$, $\varsigma_l = 0.25$, $\nu_{r,k} = 10$ and $\sigma_{r,k}$ anywhere from 0.15 to 0.50 to encourage

$\sim$30% acceptance of $\boldsymbol{\lambda}_l^*$, $\boldsymbol{\gamma}_l^*$ when $\boldsymbol{\lambda}_l \neq \mathbf{0}$ and $\varrho_{r,l}^* = 1$.

2. *Update each ϕ_m, $m = 0, 1, 2$, with MH. The proposal ϕ_m^* given the current value of ϕ_m is generated according to*

$$\phi_m^* \sim \phi_m + \sigma_{\phi_m} t(\nu_{\phi_m}) \tag{A3}$$

where σ_{ϕ_m} and ν_{ϕ_m}, $m = 0, 1, 2$ are control parameters, set in this application to $\sigma_{\phi_m} \approx 0.5$ and $\nu_{\phi_m} = 10$ to encourage $\sim$30% acceptance.

3. *Update $g_0(\boldsymbol{\omega})$ with MH. The proposal $g_0^*(\boldsymbol{\omega})$ given the current value of $g_0(\boldsymbol{\omega})$ is generated according to*

$$g_0^*(\boldsymbol{\omega}) = g_0(\boldsymbol{\omega}) + \sigma_{g_0} N(0, \boldsymbol{\Sigma}) \tag{A4}$$

where σ_{g_0} is a control parameter that is set in this application to 0.05 to encourage $\sim$30% acceptance.

4. *Update τ with MH. The proposal τ^* given the current value of τ is generated according to*

$$\tau^* = \exp\{\tau + \sigma_\tau t(\nu_\tau)\} \tag{A5}$$

where σ_τ and ν_τ are control parameters that are set in this application to 0.5 and 10, respectively, to encourage $\sim$30% acceptance.

5. *The distribution of κ given the other parameters and the data is Gamma(a_κ, b_κ), with*

$$\begin{aligned} a_\kappa &= \sum_{i=1}^{n} (N_{f,i} + N_{b,i} + \dot{h}_1) \\ b_\kappa &= \sum_{i=1}^{n} \left(\frac{20^2 \mathcal{C}_i(0)}{(20 + v_{f,i})^2} + \frac{20^2 \mathcal{C}_i(0) \exp\{-0.1185\alpha\}}{(20 + v_{b,i})^2} + \dot{h}_2 \right). \end{aligned} \tag{A6}$$

Thus, an updated value for κ is generated according to $\kappa \sim$ Gamma(a_κ, b_κ).

6. *The distribution of α_r $r = 1, 2, 3, 4$ given the other parameters and the data is Gamma$(a_{\alpha_r}, b_{\alpha_r})$, with*

$$\begin{aligned} a_{\alpha_r} &= \sum_{i=1}^{n} \left(N_{b,i} + \frac{\mu_\alpha^2}{\sigma_\alpha^2} \right) I_{\{r_i = r\}} \\ b_{\alpha_r} &= \sum_{i=1}^{n} \left(\frac{20^2 \kappa \mathcal{C}_i(0)}{(20 + v_{b,i})^2} + \frac{\mu_\alpha}{\sigma_\alpha^2} \right) I_{\{r_i = r\}}, \end{aligned} \tag{A7}$$

where $I_{\{x=y\}}$ is the indicator function, returning 1 if $x = y$, and 0 otherwise. Thus, an updated value for α_r is generated according to $\alpha_r \sim Gamma(a_{\alpha_r}, b_{\alpha_r})$.

7. *Update μ_α with MH. The proposal μ_α^* given the current value of μ_α is generated according to*

$$\mu_\alpha^* = \exp\{\mu_\alpha + \sigma_{\mu_\alpha} t(\nu_{\mu_\alpha})\} \tag{A8}$$

where σ_{μ_α} and ν_{μ_α} are control parameters that are set in this application to 0.25 and 10, respectively, to encourage ∼30% acceptance.

8. *Update σ_α^2 with MH. The proposal $\sigma_\alpha^{2^*}$ given the current value of σ_α^2 is generated according to*

$$\sigma_\alpha^{2^*} = \exp\{\sigma_\alpha^2 + \sigma_{\sigma_\alpha^2} t(\nu_{\sigma_\alpha^2})\} \tag{A9}$$

where $\sigma_{\sigma_\alpha^2}$ and $\nu_{\sigma_\alpha^2}$ are control parameters that are set in this application to 1.5 and 10, respectively, to encourage ∼30% acceptance.